Variable length trajectory compressible hybrid Monte Carlo

Akihiko Nishimura[1, a)] and David Dunson[2]
[1)] *Department of Mathematics, Duke University, Durham, North Carolina, 27708, USA*
[2)] *Department of Statistical Science, Duke University, Durham, North Carolina, 27708, USA*

(Dated: 5 April 2016)

Hybrid Monte Carlo (HMC) generates samples from a prescribed probability distribution in a configuration space by simulating Hamiltonian dynamics, followed by the Metropolis (-Hastings) acceptance/rejection step. Compressible HMC (CHMC) generalizes HMC to a situation in which the dynamics is reversible but not necessarily Hamiltonian. This article presents a framework to further extend the algorithm. Within the existing framework, each trajectory of the dynamics must be integrated for the same amount of (random) time to generate a valid Metropolis proposal. Our generalized acceptance/rejection mechanism allows a more deliberate choice of the integration time for each trajectory. The proposed algorithm in particular enables an effective application of variable step size integrators to HMC-type sampling algorithms based on reversible dynamics. The potential of our framework is further demonstrated by another extension of HMC which reduces the wasted computations due to unstable numerical approximations and corresponding rejected proposals.

arXiv:1604.00889v1 [physics.comp-ph] 4 Apr 2016

I. INTRODUCTION

A study of molecular systems often relies on generating random variables from a prescribed (unnormalized) probability distribution $\rho(\boldsymbol{q}) \propto \exp(-U(\boldsymbol{q}))$ on the configuration space. Markov chain Monte Carlo (MCMC) generates samples from a target distribution by constructing a Markov chain whose stationary distribution coincides with the target distribution. Such a Markov chain can be realized by building a transition rule that satisfies the detailed balance condition. MCMC based on the Metropolis (-Hastings) algorithm[1] is a general sampling approach widely used in computational physical science as well as in Bayesian statistics and machine learning. Many such algorithms are inefficient, producing highly correlated samples, and require a large number of iterations to adequately characterize the target distribution[2–4].

Hybrid Monte Carlo[5] (HMC) constructs the proposal distribution for the Metropolis algorithm by simulating molecular dynamics (MD), a procedure that can greatly reduce the correlation among successive MCMC samples. More precisely, HMC augments the state space by introducing a *momentum* variable $\boldsymbol{p}$; the original variable $\boldsymbol{q}$ is often referred to as *position* variable in the HMC framework. In this augmented state space $(\boldsymbol{q}, \boldsymbol{p})$, the proposal distribution for Metropolis is constructed by solving an ordinary differential equation (ODE) corresponding to Newton's equations of motion with respect to the potential energy $U(\boldsymbol{q})$. There are applications in which (partial) analytical solutions to an ODE can be exploited[6–8], but in general ODEs are discretized and integrated numerically.

Within the original HMC framework, an integrator for simulating MD must be reversible and volume-preserving to produce a valid Metropolis proposal[9]. In fact, the volume-preserving property can be relaxed by including a Jacobian factor in the calculation of the Metropolis acceptance probability.[10,11] Under this generalization of HMC, any reversible (discrete) dynamics / bijective map can be applied to generate a proposal state. This algorithm is formalized as *compressible HMC* (CHMC) in Ref. 12. A generalization of HMC known as *Riemannian manifold HMC* (RMHMC)[13] also falls within the framework of CHMC.

This article presents an algorithm to relax another condition required by (compressible) HMC. Given a reversible map $\boldsymbol{F}$ and state $(\boldsymbol{q}, \boldsymbol{p})$, CHMC proposes the next state by applying the map n times, where the number of steps n can be drawn randomly at each iteration. Though often not stated explicitly, the detailed balance requires the number of steps to be determined independently of the trajectory $\{(\boldsymbol{q}, \boldsymbol{p}), \boldsymbol{F}(\boldsymbol{q}, \boldsymbol{p}), \boldsymbol{F}^2(\boldsymbol{q}, \boldsymbol{p}), \ldots\}$. As we will show in Section III, this constraint can prevent realizing the full potential of MCMC algorithms based on reversible dynamics.

Our algorithm generalizes the acceptance-rejection mechanism behind CHMC to allow the number of steps to depend on each trajectory of the dynamics while preserving the detailed balance. The number of numerical integration steps taken in simulating a trajectory of HMC is commonly referred to as the "path length" of a trajectory in the statistics literature. We therefore call our algorithm *variable length trajectory CHMC* (VLT-CHMC). It should be mentioned that the No-U-Turn-Sampler (NUTS) is another variant of HMC that allows the path lengths to vary from one trajectory to another.[14] However, the motivation behind NUTS is to spare a user the trouble of manually tuning the number of steps, and NUTS in general performs no better than HMC with well-chosen path lengths.[14,15] On the other hand VLT-CHMC can improve the performance of CHMC in a more fundamental and significant way. In particular, VLT-CHMC enables an effective application of reversible variable step size integrators to HMC-type sampling algorithms based on reversible dynamics.

[a)] Electronic mail: an88@duke.edu.

The rest of the paper is organized as follows. Section II reviews the main ideas behind CHMC and provides an example in which the compressible dynamics arises from the use of non-traditional integrators in HMC settings. Such integrators have proven to be more efficient than the commonly used volume-preserving integrators in various applications. The example also serves to introduce the notations and concepts needed in the next section, where VLT-CHMC is motivated as a method to effectively apply variable step size integrators in HMC settings. The presentation is self-contained, but some familiarity with HMC is assumed. VLT-CHMC is developed in Section III. Section III A explains how the existing framework limits the utility of variable step size integrators to sampling algorithms. The key observation in addressing this issue leads to a special case of VLT-CHMC. More general construction of VLT-CHMC is provided in Section III B. Section III C presents another use case of VLT-CHMC, where HMC is modified to reduce the wasted computation due to unstable numerical approximations and corresponding rejected proposals. The simulation results are shown in Section IV to demonstrate the potential gains from the framework of VLT-CHMC.

II. REVIEW OF COMPRESSIBLE HMC

A. Basic Theory

To keep the description of CHMC and the subsequent development of VLT-CHMC more intuitive, the version of CHMC described here is slightly less general than the one in Ref. 12. It is straightforward to extend the variable length trajectory algorithm of Section III to the general settings.

A bijective map $\boldsymbol{F}$ is said to be *reversible* if

$$\boldsymbol{F}^{-1} = \boldsymbol{R} \circ \boldsymbol{F} \circ \boldsymbol{R} \tag{1}$$

or equivalently $(\boldsymbol{R} \circ \boldsymbol{F})^{-1} = \boldsymbol{R} \circ \boldsymbol{F}$ for an involution $\boldsymbol{R}$ (i.e. $\boldsymbol{R} \circ \boldsymbol{R} = \mathrm{id}$). Note that the reversiblity of $\boldsymbol{F}$ implies that of $\boldsymbol{F}^n$ for any n. Let $\boldsymbol{DF}^n$ denote the Jacobian matrix of $\boldsymbol{F}^n$ and $|\boldsymbol{DF}^n|$ its determinant. Given a state $\boldsymbol{z}$ and integer n, CHMC proposes the state $\boldsymbol{z}^* = \boldsymbol{R} \circ \boldsymbol{F}^n(\boldsymbol{z})$ and accepts or rejects the proposal with probability

$$\min\left\{1, \frac{\rho(\boldsymbol{z}^*)|\boldsymbol{DF}^n(\boldsymbol{z})|}{\rho(\boldsymbol{z})}\right\} \tag{2}$$

To see that this transition rule satisfies the detailed balance with respect to $\rho(\cdot)$, consider a small neighborhood B around $\boldsymbol{z}$ and $B^* = \boldsymbol{R} \circ \boldsymbol{F}^n(B)$ around $\boldsymbol{z}^*$, so that $\boldsymbol{R} \circ \boldsymbol{F}^n(B^*) = B$. The proposal move sends the probability mass

$$\int_B \rho(\boldsymbol{z}')\mathrm{d}\boldsymbol{z}' \approx \rho(\boldsymbol{z})\mathrm{vol}(B)$$

from B to B^*. On the other hand, the mass sent from B^* to B by the proposal move can be seen to be

$$\begin{aligned}\int_{B^*} \rho(\boldsymbol{z}')\mathrm{d}\boldsymbol{z}' &= \int_B \rho(\boldsymbol{z}')|\boldsymbol{D}(\boldsymbol{R} \circ \boldsymbol{F}^n)(\boldsymbol{z}')|\mathrm{d}\boldsymbol{z}' \\ &\approx \rho(\boldsymbol{z}^*)|\boldsymbol{DF}^n(\boldsymbol{z})|\,\mathrm{vol}(B)\end{aligned}$$

by the change of variable formula and the fact $|\boldsymbol{R}| = 1$. The acceptance and rejection step of CHMC amounts to rejecting the fraction of move by the ratio of the probability fluxes and thus imposes the detailed balance.

The above transition rule preserves the target density $\rho(\cdot)$ for any n, so in practice the number of steps can be drawn randomly at each iteration of CHMC. The steps of CHMC are summarized in Algorithm 1 below, where the distribution $p(\cdot)$ for the number of steps is a tuning parameter a user must specify. The use of a deterministic map as a proposal distribution does not yield an ergodic Markov chain, and therefore such a transition rule must be alternated with another transition rule that preserves the target density $\rho(\cdot)$, as done in Step 1 of the algorithm. We do not concern ourselves here with how to choose such a random move since the choice depends critically on the particular form of $\rho(\cdot)$.

Algorithm 1 (Compressible HMC)**.** With a prespecified probability mass function $p(\cdot)$ on $\mathbb{Z}^+$, CHMC generates a Markov chain $\{\boldsymbol{z}^{(m)}\}_m$ with the following transition rule $\boldsymbol{z}^{(m)} \to \boldsymbol{z}^{(m+1)}$:

1. Make a random change $\boldsymbol{z}^{(m)} \to \boldsymbol{z}$ that preserves the target density $\rho(\cdot)$.
2. Sample $n \sim p(\cdot)$ and propose the state $\boldsymbol{z}^* = \boldsymbol{R} \circ \boldsymbol{F}^n(\boldsymbol{z})$.
3. Let $\boldsymbol{z}^{(m+1)} = \boldsymbol{z}^*$ with probability
$$\min\left\{1, \frac{\rho(\boldsymbol{z}^*)|\boldsymbol{DF}^n(\boldsymbol{z})|}{\rho(\boldsymbol{z})}\right\}$$
Otherwise, let $\boldsymbol{z}^{(m+1)} = \boldsymbol{z}$.

B. Example: (Riemann manifold) HMC with non-volume-preserving integrators

HMC and its extension Riemann manifold HMC (RMHMC) construct a reversible and volume-preserving bijective map by numerically approximating Hamiltonian dynamics. To this end, they require a geometric integrator that preserves the reversibility and volume-preservation property of Hamiltonian dynamics. Under the CHMC framework, however, Hamiltonian dynamics can be approximated using a wider range of integration techniques.

In order to sample from a probability density of interest $\rho_0(\boldsymbol{q}) \propto \exp(-U(\boldsymbol{q}))$ in $\mathbb{R}^d$, RMHMC introduces an auxiliary variable $\boldsymbol{p} \in \mathbb{R}^d$ whose distribution is defined conditionally as $\boldsymbol{p}|\boldsymbol{q} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{M}(\boldsymbol{q}))$ for a family of positive definite matrices known as *mass tensors*

$\{\boldsymbol{M}(\boldsymbol{q})\}_{\boldsymbol{q}}$.[9,13,16] The joint density $\rho(\boldsymbol{q},\boldsymbol{p})$ in the phase space then is given as $\rho(\boldsymbol{q},\boldsymbol{p}) \propto \exp(-H(\boldsymbol{q},\boldsymbol{p}))$ where the *Hamiltonian* $H(\boldsymbol{q},\boldsymbol{p})$ is given by

$$H(\boldsymbol{q},\boldsymbol{p}) = U(\boldsymbol{q}) + \frac{1}{2}\boldsymbol{p}^T\boldsymbol{M}(\boldsymbol{q})^{-1}\boldsymbol{p} + \frac{1}{2}\log|\boldsymbol{M}(\boldsymbol{q})| \quad (3)$$

The proposal is generated by approximating the solution to *Hamilton's equations*:

$$\frac{\mathrm{d}\boldsymbol{q}}{\mathrm{d}t} = \nabla_{\boldsymbol{p}}H(\boldsymbol{q},\boldsymbol{p}),\ \frac{\mathrm{d}\boldsymbol{p}}{\mathrm{d}t} = -\nabla_{\boldsymbol{q}}H(\boldsymbol{q},\boldsymbol{p}) \quad (4)$$

For the Hamiltonian (3), the solution operator of (4) is reversible with respect to a momentum flip operator $\boldsymbol{R}(\boldsymbol{q},\boldsymbol{p}) = (\boldsymbol{q},-\boldsymbol{p})$. Solving (4) using a reversible integrator with a constant step size Δt yields a reversible map $\boldsymbol{F}_{\Delta t}$ so that

$$\boldsymbol{F}^n_{\Delta t}(\boldsymbol{q}_0,\boldsymbol{p}_0) \approx (\boldsymbol{q}(n\Delta t),\boldsymbol{p}(n\Delta t)) \quad (5)$$

where $\{(\boldsymbol{q}(t),\boldsymbol{p}(t))\}_t$ denotes the exact solution with the initial condition $(\boldsymbol{q}_0,\boldsymbol{p}_0)$. In other words, $\boldsymbol{F}_{\Delta t}$ approximates the solution operator $\boldsymbol{\Phi}_{\Delta t}$ of (4) defined through the relation

$$\frac{\mathrm{d}\boldsymbol{\Phi}_t}{\mathrm{d}t} = \big((\nabla_{\boldsymbol{p}}H)\circ\boldsymbol{\Phi}_t, -(\nabla_{\boldsymbol{q}}H)\circ\boldsymbol{\Phi}_t\big) \quad (6)$$

for all t. If the reversible map $\boldsymbol{F}_{\Delta t}$ is further required to be volume preserving, then we have $|\boldsymbol{D}\boldsymbol{F}^n_{\Delta t}| = 1$ and the Jacobian factor drops from (2), recovering HMC and RMHMC algorithms of Ref. 5 and 13. In some applications however, non-volume-preserving approximations of (4) have been shown to offer substantial gains in computational efficiency.[11,12]

For example, Lan et. al.[11] considers the ODE corresponding to (4) in terms of reparametrization $(\boldsymbol{q},\mathbf{v}) = (\boldsymbol{q},\boldsymbol{M}(\boldsymbol{q})^{-1}\boldsymbol{p})$. The reparametrized ODE admits semi-explicit and explicit reversible approximations, requiring fewer or no fixed point iterations compared to the Störmer-Verlet integrator typically employed in RMHMC. The proposal move using a simulated trajectory is alternated with sampling $\mathbf{v}$ from its conditional density $\mathbf{v}|\boldsymbol{q} \sim \mathcal{N}(\mathbf{0},\boldsymbol{M}(\boldsymbol{q})^{-1})$, a random move corresponding to Step 1 in Algorithm 1. The CHMC algorithm based on the semi-explicit and explicit integrator are found to significantly outperform RMHMC based on the Störmer-Verlet integrator over a range of examples.

III. VARIABLE LENGTH TRAJECTORY CHMC

Variable length trajectory CHMC (VLT-CHMC) is most naturally motivated as a method to effectively apply variable step size integrators in RMHMC settings. For this reason, we first develop this special case of VLT-CHMC in Section III A. A more general theory is developed in Section III B. Section III C illustrates the use and potential benefits of the general VLT-CHMC algorithm through another example.

A. Special case of VLT-CHMC

1. Motivation: RMHMC with variable step size integrators and limitations of CHMC

In Section II B, we discussed how CHMC allows us to approximate Hamiltonian dynamics with non-volume-preserving integrators and still generate a valid Metropolis proposal. We in particular considered the use of a reversible integrator with a constant step size. A wider range of reversible integration techniques for Hamiltonian systems are available in the literature, however, including a number of variable step size integrators.[17–20] In theory, a variable step size integrator similarly produces a valid CHMC proposal as long as the integrator is reversible. However, the use of such an integrator under the existing CHMC framework generally leads to an algorithm with suboptimal sampling efficiency, for the reasons we describe now.

Each step of a variable step size integrator approximates the evolution $(\boldsymbol{q}(t_n),\boldsymbol{p}(t_n)) \to (\boldsymbol{q}(t_n+\Delta t_n),\boldsymbol{p}(t_n+\Delta t_n))$ where the step size Δt_n depends on the current state $(\boldsymbol{q}(t_n),\boldsymbol{p}(t_n))$ through a *step size controller* $g(\boldsymbol{q},\boldsymbol{p})$. The simplest choice of step size would be $\Delta t_n = g(\boldsymbol{q}(t_n),\boldsymbol{p}(t_n))\Delta s$, but the reversibility requires a slightly more sophisticated relationship and the condition $g(\boldsymbol{q},\boldsymbol{p}) = g(\boldsymbol{q},-\boldsymbol{p})$ (see Section III A 2). Most importantly for our discussion, a variable step size scheme is equivalent to approximating the following *time-rescaled* Hamiltonian dynamics in a new time scale $\mathrm{d}s = g(\boldsymbol{q},\boldsymbol{p})^{-1}\mathrm{d}t$ with a constant step size Δs:

$$\frac{\mathrm{d}\boldsymbol{q}}{\mathrm{d}s} = g(\boldsymbol{q},\boldsymbol{p})\nabla_{\boldsymbol{p}}H(\boldsymbol{q},\boldsymbol{p}),\ \frac{\mathrm{d}\boldsymbol{p}}{\mathrm{d}s} = -g(\boldsymbol{q},\boldsymbol{p})\nabla_{\boldsymbol{q}}H(\boldsymbol{q},\boldsymbol{p}) \quad (7)$$

In other words, a reversible variable step size approximation of (4) yields a reversible map $\boldsymbol{F}_{\Delta s}$ such that

$$\boldsymbol{F}^n_{\Delta s}(\boldsymbol{q}_0,\boldsymbol{p}_0) \approx (\boldsymbol{q}(n\Delta s),\boldsymbol{p}(n\Delta s)) \quad (8)$$

where $\{\boldsymbol{q}(s),\boldsymbol{p}(s)\}_s$ is the solution to the time-rescaled dynamics (7) with the initial condition $(\boldsymbol{q}_0,\boldsymbol{p}_0)$.

This implicit time-rescaling behind variable step size integration causes trouble for CHMC. The utility of Hamiltonian dynamics (4) as a proposal generation mechanism stems from the fact that $\rho(\boldsymbol{q},\boldsymbol{p}) \propto \exp(-H(\boldsymbol{q},\boldsymbol{p}))$ is the *invariant distribution* of the dynamics i.e. if $(\boldsymbol{q}_0,\boldsymbol{p}_0)$ has the distribution $\rho(\boldsymbol{q},\boldsymbol{p}) \propto \exp(-H(\boldsymbol{q},\boldsymbol{p}))$, then $\boldsymbol{\Phi}_t(\boldsymbol{q}_0,\boldsymbol{p}_0)$ also has the same distribution $\rho(\cdot)$ for all t. As a consequence, the proposal generated by an approximate solution $(\boldsymbol{q}^*,\boldsymbol{p}^*) = \boldsymbol{F}^n_{\Delta t}(\boldsymbol{q}_0,\boldsymbol{p}_0)$ as in (5) can be accepted with probability 1 in the limit $\Delta t \to 0$ and $n\Delta t \to t'$. On the other hand, the time-rescaled dynamics (7) in general does not preserve the target density $\rho(\boldsymbol{q},\boldsymbol{p})$, and the proposal generated by the approximate solution $(\boldsymbol{q}^*,\boldsymbol{p}^*) = \boldsymbol{F}^n_{\Delta s}(\boldsymbol{q}_0,\boldsymbol{p}_0)$ may not be accepted with high probability even in the limit $\Delta s \to 0$ and $n\Delta s \to s'$. In fact, the acceptance probability of the CHMC proposal

in the limit is given by:

$$\min\left\{1, \frac{g(\boldsymbol{q}(s'), \boldsymbol{p}(s'))}{g(\boldsymbol{q}_0, \boldsymbol{p}_0)}\right\} \tag{9}$$

where $\{\boldsymbol{q}(s), \boldsymbol{p}(s)\}_s$ denotes the solution to (7) with the initial condition $(\boldsymbol{q}_0, \boldsymbol{p}_0)$. The derivation is given in Appendix A.

2. Algorithm: variable length trajectory scheme for time-rescaled dynamics

In order to address the issue caused by the implicit time-rescaling associated with variable step size integrators, VLT-CHMC approximates the dynamics in the original time scale as follows. Fix the initial condition $(\boldsymbol{q}_0, \boldsymbol{p}_0)$ and denote $(\boldsymbol{q}_i, \boldsymbol{p}_i) = \boldsymbol{F}^i_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0)$ where $\boldsymbol{F}_{\Delta s}$ approximates the dynamics in the time scale s as in (8). The evolution $(\boldsymbol{q}_0, \boldsymbol{p}_0) \to (\boldsymbol{q}(t), \boldsymbol{p}(t))$ in the original time scale can be approximated by taking the trajectory dependent number of steps $N(\boldsymbol{q}_0, \boldsymbol{p}_0) = N(t, \boldsymbol{q}_0, \boldsymbol{p}_0)$ defined as

$$N(\boldsymbol{q}_0, \boldsymbol{p}_0) = \min\left\{n : \sum_{i=1}^{n} \frac{\Delta s}{2}\left(g(\boldsymbol{q}_{i-1}, \boldsymbol{p}_{i-1}) + g(\boldsymbol{q}_i, \boldsymbol{p}_i)\right) > t\right\} \tag{10}$$

Now we consider the map $\boldsymbol{F}^N_{\Delta s}$ defined as

$$\boldsymbol{F}^N_{\Delta s}(\boldsymbol{q}, \boldsymbol{p}) = \boldsymbol{F}^{N(\boldsymbol{q},\boldsymbol{p})}_{\Delta s}(\boldsymbol{q}, \boldsymbol{p}) \tag{11}$$

which approximates the solution operator $\boldsymbol{\Phi}_t$ as defined in (6). The map however cannot be used directly to generate a proposal because in general it is neither reversible or even bijective. The map would be reversible if $N(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*) = N(\boldsymbol{q}_0, \boldsymbol{p}_0)$ where $(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*) = \boldsymbol{R} \circ \boldsymbol{F}^N_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0)$, but (10) only implies $N(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*) \le N(\boldsymbol{q}_0, \boldsymbol{p}_0)$. For example when $g(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*) \gg g(\boldsymbol{q}_0, \boldsymbol{p}_0)$, the simulated time along the reverse trajectory $\left\{(\boldsymbol{q}_i^*, \boldsymbol{p}_i^*) = \boldsymbol{F}^i_{\Delta s}(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*)\right\}_{i=0}^n$

$$\sum_{i=1}^{n} \frac{\Delta s}{2}\big(g(\boldsymbol{q}_{i-1}^*, \boldsymbol{p}_{i-1}^*) + g(\boldsymbol{q}_i^*, \boldsymbol{p}_i^*)\big)$$

will likely reach the threshold t before $n = N(\boldsymbol{q}_0, \boldsymbol{p}_0)$ steps.

The key observation behind VLT-CHMC is that we can nonetheless construct collections of states S and S^* containing $(\boldsymbol{q}_0, \boldsymbol{p}_0)$ and $(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*)$ such that

$$\begin{aligned} \boldsymbol{R} \circ \boldsymbol{F}^N_{\Delta s}(S) \subset S^* \quad \text{and} \quad \boldsymbol{R} \circ \boldsymbol{F}^N_{\Delta s}(S^*) \subset S \\ \boldsymbol{R} \circ \boldsymbol{F}^N_{\Delta s}(S^c) \subset (S^*)^c \quad \text{and} \quad \boldsymbol{R} \circ \boldsymbol{F}^N_{\Delta s}\left((S^*)^c\right) \subset S^c \end{aligned} \tag{12}$$

The existence of such sets S and S^* is a property of the map $\boldsymbol{F}^N_{\Delta s}$ and generalizes the notion of reversibility (1). The set S is essentially the pre-image of $\{(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*)\}$ under $\boldsymbol{R} \circ \boldsymbol{F}^N_{\Delta s}$ and can be constructed by defining $S = \{(\boldsymbol{q}_{-\ell}, \boldsymbol{p}_{-\ell}), \ldots, (\boldsymbol{q}_r, \boldsymbol{p}_r)\}$ by choosing $\ell, r \ge 0$ such that

$$\begin{aligned} \ell &= \max\left\{j \ge 0 : \boldsymbol{F}^N_{\Delta s}(\boldsymbol{q}_{-j}, \boldsymbol{p}_{-j}) = \boldsymbol{F}^N_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0)\right\} \\ r &= \max\left\{j \ge 0 : \boldsymbol{F}^N_{\Delta s}(\boldsymbol{q}_j, \boldsymbol{p}_j) = \boldsymbol{F}^N_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0)\right\} \end{aligned} \tag{13}$$

Algorithmically, ℓ and r can be found by solving the dynamics backward and forward from $(\boldsymbol{q}_0, \boldsymbol{p}_0)$ using the equivalent definitions below:

$$\begin{aligned} \ell &= \max\left\{j \ge 0 : \sum_{i=-j}^{N(t,\boldsymbol{q}_0,\boldsymbol{p}_0)-1} \Delta t_i < t\right\} \\ r &= \max\left\{j \ge 0 : \sum_{i=j}^{N(t,\boldsymbol{q}_0,\boldsymbol{p}_0)} \Delta t_i > t\right\} \\ &\text{where } \Delta t_i = \frac{\Delta s}{2}\left(g(\boldsymbol{q}_{i-1}, \boldsymbol{p}_{i-1}) + g(\boldsymbol{q}_i, \boldsymbol{p}_i)\right) \end{aligned} \tag{14}$$

The set S^* is the pre-image of $\{(\boldsymbol{q}_r, \boldsymbol{p}_r)\}$ under $\boldsymbol{R} \circ \boldsymbol{F}^N_{\Delta s}$ and can analogously be constructed. Denoting $(\boldsymbol{q}_i^*, \boldsymbol{p}_i^*) = \boldsymbol{F}^i_{\Delta s}(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*)$, let $S^* = \left\{(\boldsymbol{q}^*_{-\ell^*}, \boldsymbol{p}^*_{-\ell^*}), \ldots, (\boldsymbol{q}^*_{r^*}, \boldsymbol{p}^*_{r^*})\right\}$ where $\ell^*, r^* \ge 0$ is defined as

$$\begin{aligned} \ell^* &= \max\left\{j \ge 0 : \boldsymbol{F}^N_{\Delta s}(\boldsymbol{q}^*_{-j}, \boldsymbol{p}^*_{-j}) = \boldsymbol{F}^N_{\Delta s}(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*)\right\} \\ r^* &= \max\left\{j \ge 0 : \boldsymbol{F}^N_{\Delta s}(\boldsymbol{q}_j^*, \boldsymbol{p}_j^*) = \boldsymbol{F}^N_{\Delta s}(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*)\right\} \end{aligned} \tag{15}$$

It is shown in Appendix B that the above definition actually implies $r^* = 0$. The proof of (12) and of other facts regarding S and S^* are also given in Appendix B.

Having constructed the sets S and S^* with the property (12), VLT-CHMC imposes the detailed balance by rejecting a fraction of moves between S and S^* as described in Algorithm 2 below.

Algorithm 2 (VLT-CHMC)**.** Given a reversible map $\boldsymbol{F}_{\Delta s}$ as in (8) and a trajectory length function N as in (10), VLT-CHMC generates a Markov chain $\{(\boldsymbol{q}^{(m)}, \boldsymbol{p}^{(m)})\}_m$ with the following transition rule $(\boldsymbol{q}^{(m)}, \boldsymbol{p}^{(m)}) \to (\boldsymbol{q}^{(m+1)}, \boldsymbol{p}^{(m+1)})$:

1. Sample $\boldsymbol{p}_0$ from the conditional density $\boldsymbol{p}|\boldsymbol{q}^{(m)}$ and set $\boldsymbol{q}_0 = \boldsymbol{q}^{(m)}$.
2. Find the indices ℓ, r, ℓ^*, r^* as in (13) and (15) by simulating the dynamics forward and backward from $(\boldsymbol{q}_0, \boldsymbol{p}_0)$ and $\boldsymbol{R} \circ \boldsymbol{F}^N_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0)$. Then set
$$\begin{aligned} S &= \left\{\boldsymbol{F}^{-\ell}_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0), \ldots, \boldsymbol{F}^{r}_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0)\right\} \\ S^* &= \Big\{\boldsymbol{R} \circ \boldsymbol{F}^{N_0 - r^*}_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0), \\ &\qquad \ldots, \boldsymbol{R} \circ \boldsymbol{F}^{N_0 + \ell^*}_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0)\Big\} \end{aligned}$$
where $N_0 = N(\boldsymbol{q}_0, \boldsymbol{p}_0)$.
3. Propose the transition from S to S^* with the acceptance probability which is the smaller of 1 and
$$\frac{\sum_{j=-r*}^{\ell^*} \rho\left(\boldsymbol{R} \circ \boldsymbol{F}^{N_0+j}_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0)\right)\left|\boldsymbol{D}\boldsymbol{F}^{N_0+j}_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0)\right|}{\sum_{i=-\ell}^{r} \rho\left(\boldsymbol{F}^i_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0)\right)\left|\boldsymbol{D}\boldsymbol{F}^{i}_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0)\right|} \tag{16}$$

4. If the transition in Step 3 is accepted, choose a state $\boldsymbol{R}\circ\boldsymbol{F}_{\Delta s}^{N_0+j}(\boldsymbol{q}_0,\boldsymbol{p}_0)$ from S^* with the probability proportional to

$$\rho\left(\boldsymbol{R}\circ\boldsymbol{F}_{\Delta s}^{N_0+j}(\boldsymbol{q}_0,\boldsymbol{p}_0)\right)\left|\boldsymbol{D}\boldsymbol{F}_{\Delta s}^{N_0+j}(\boldsymbol{q}_0,\boldsymbol{p}_0)\right| \tag{17}$$

and set $(\boldsymbol{q}^{(m+1)},\boldsymbol{p}^{(m+1)}) = \boldsymbol{R}\circ\boldsymbol{F}_{\Delta s}^{N_0+j}(\boldsymbol{q}_0,\boldsymbol{p}_0)$. Otherwise, choose a state $\boldsymbol{F}_{\Delta s}^{i}(\boldsymbol{q}_0,\boldsymbol{p}_0)$ from S with the probability proportional to

$$\rho\left(\boldsymbol{F}_{\Delta s}^{i}(\boldsymbol{q}_0,\boldsymbol{p}_0)\right)\left|\boldsymbol{D}\boldsymbol{F}_{\Delta s}^{i}(\boldsymbol{q}_0,\boldsymbol{p}_0)\right| \tag{18}$$

and set $(\boldsymbol{q}^{(m+1)},\boldsymbol{p}^{(m+1)}) = \boldsymbol{F}_{\Delta s}^{i}(\boldsymbol{q}_0,\boldsymbol{p}_0)$.

3. Theory: VLT-CHMC and detailed-balance condition

Too see how VLT-CHMC achieves the detailed balance, consider a small neighborhood B_0 around $(\boldsymbol{q}_0,\boldsymbol{p}_0)$. The total probability in the neighborhood $B = \cup_{i=-\ell}^{r}\boldsymbol{F}_{\Delta s}^{i}(B_0)$ of S is

$$\begin{aligned}&\int_B \rho(\boldsymbol{q},\boldsymbol{p})\,\mathrm{d}\boldsymbol{q}\,\mathrm{d}\boldsymbol{p}\\ &\quad\approx \sum_{i=-\ell}^{r}\rho\left(\boldsymbol{F}_{\Delta s}^{i}(\boldsymbol{q}_0,\boldsymbol{p}_0)\right)\left|\boldsymbol{D}\boldsymbol{F}_{\Delta s}^{i}(\boldsymbol{q}_0,\boldsymbol{p}_0)\right|\left|B_0\right|\end{aligned} \tag{19}$$

assuming that B_0 is small enough that $\boldsymbol{F}_{\Delta s}^{i}(B_0)$'s are disjoint. Similarly, the total probability in the neighborhood $B^* = \cup_{j=-r^*}^{\ell^*}\boldsymbol{R}\circ\boldsymbol{F}_{\Delta s}^{N_0+j}(B_0)$ of S^* is

$$\begin{aligned}&\int_{B^*} \rho(\boldsymbol{q},\boldsymbol{p})\,\mathrm{d}\boldsymbol{q}\,\mathrm{d}\boldsymbol{p}\\ &\quad\approx \sum_{j=-r^*}^{\ell^*}\rho\left(\boldsymbol{R}\circ\boldsymbol{F}_{\Delta s}^{N_0+j}(\boldsymbol{q}_0,\boldsymbol{p}_0)\right)\left|\boldsymbol{D}\boldsymbol{F}_{\Delta s}^{N_0+j}(\boldsymbol{q}_0,\boldsymbol{p}_0)\right|\left|B_0\right|\end{aligned} \tag{20}$$

Comparing the acceptance probability (16) with the probability fluxes (19) and (20), one can see that the acceptance-rejection procedure of Step 3 controls the probability fluxes appropriately to achieve the detailed balance between the neighborhoods B and B^*. Step 4 then imposes the detailed balance within B and B^* by sampling a state according to the relative amount of probability in the individual components $\{\boldsymbol{F}_{\Delta s}^{i}(B_0)\}_{i=-\ell}^{r}$ of B and $\{\boldsymbol{R}\circ\boldsymbol{F}_{\Delta s}^{N_0+j}(B_0)\}_{j=-r^*}^{\ell^*}$ of B^*.

4. Theoretical efficiency: improvement over CHMC

Throughout Section III A we considered the compressible dynamics (7) arising from a variable step size integration of Hamiltonian dynamics. In this specific setting with the trajectory length function N as defined in (10), VLT-CHMC is guaranteed to have a high average acceptance probability. In fact, in the limit $\Delta s\to 0$ with t fixed, the acceptance probability (16) of a VLT-CHMC proposal from $(\boldsymbol{q}_0,\boldsymbol{p}_0)$ converges to a value bounded below by

$$\frac{g(\boldsymbol{\Phi}_t(\boldsymbol{q}_0,\boldsymbol{p}_0))}{g(\boldsymbol{q}_0,\boldsymbol{p}_0)}\left\lfloor\frac{g(\boldsymbol{q}_0,\boldsymbol{p}_0)}{g(\boldsymbol{\Phi}_t(\boldsymbol{q}_0,\boldsymbol{p}_0))}\right\rfloor \tag{21}$$

when $g(\boldsymbol{\Phi}_t(\boldsymbol{q}_0,\boldsymbol{p}_0)) < g(\boldsymbol{q}_0,\boldsymbol{p}_0)$. In case $g(\boldsymbol{\Phi}_t(\boldsymbol{q}_0,\boldsymbol{p}_0)) > g(\boldsymbol{q}_0,\boldsymbol{p}_0)$, a similar lower bound holds for the proposal from $\boldsymbol{R}\circ\boldsymbol{\Phi}_t(\boldsymbol{q}_0,\boldsymbol{p}_0)$. Note that the quantity (21) is always larger than $1/2$ and it tends to 1 as the ratio $g(\boldsymbol{\Phi}_t(\boldsymbol{q}_0,\boldsymbol{p}_0))/g(\boldsymbol{q}_0,\boldsymbol{p}_0)$ increases, in contrast with the acceptance probability (9) of CHMC. More precise results on the acceptance probability of a VLT-CHMC proposal are derived in Appendix A.

Of course, the acceptance rate of a proposal distribution is not the only factor determining the efficiency of an MCMC algorithm. Nonetheless, the theoretical result above highlights an advantage VLT-CHMC has over the usual CHMC. The bottom line is that VLT-CHMC proposals approximate the original dynamic (4) while CHMC proposals approximate the time-rescaled dynamics (7). Therefore, VLT-CHMC will generally outperform CHMC whenever the exact solution of the original dynamics constitutes an efficient Markov chain propagator as is typically the case in RMHMC applications.[13,21] This is substantiated by our simulation study in Section IV.

B. General VLT-CHMC

The key step in Algorithm 2 is the construction of the sets S and S^* with the property (12). More generally, the detailed balance can be imposed by the same type of acceptance-rejection mechanism whenever the phase space can be partitioned into a collection of pairs S and S^* such that the set $S\cup S^*$ and $(S\cup S^*)^c$ is closed under a (deterministic) transition rule. Conceivably, a wide range of algorithms can be devised under this general condition. In this section we present one systematic way to generalize the framework of Section III A.

Consider a generic reversible map $\boldsymbol{F}$ on a state space $\boldsymbol{z}$ and associated involution $\boldsymbol{R}$. Fix $\boldsymbol{z}_0$ and denote $\boldsymbol{z}_i = \boldsymbol{F}^i(\boldsymbol{z}_0)$. Choose a *trajectory termination criteria*, or more precisely boolean valued functions $b_n(\boldsymbol{z}_0,\dots,\boldsymbol{z}_n)\in\{0,1\}$, with the following property

$$b_n(\boldsymbol{z}_0,\dots,\boldsymbol{z}_n) = b_n(\boldsymbol{R}(\boldsymbol{z}_n),\dots,\boldsymbol{R}(\boldsymbol{z}_0)) \tag{22}$$

as well as the property

$$b_n(\boldsymbol{z}_0,\dots,\boldsymbol{z}_n) = 1 \;\text{ only if }\; b_{n-i}(\boldsymbol{z}_i,\dots,\boldsymbol{z}_n) = 1 \tag{23}$$

for any $i>0$. These properties are satisfied, for example, by a termination criteria $\sum_{i=1}^{n} a(\boldsymbol{z}_i) + a(\boldsymbol{z}_{i-1}) > c$ for a scalar function $a(\boldsymbol{z})\ge 0$. Define a corresponding trajectory length function $N(\boldsymbol{z}_0)$ as

$$\begin{aligned}&N(\boldsymbol{z}_0) = \min\{N'(\boldsymbol{z}_0), N_{\max}\}\\ &\qquad\text{for }\; N'(\boldsymbol{z}_0) = \min\left\{n : b_n(\boldsymbol{z}_0,\dots,\boldsymbol{z}_n) = 1\right\}\end{aligned} \tag{24}$$

With the reversible map $\boldsymbol{F}_{\Delta s}$ and trajectory length function N of (10) replaced by the generic ones as above, Algorithm 2 remains a valid MCMC scheme. This is because the justification of the algorithm (in Appendix B) only require a trajectory length function N to satisfy the *short return* condition

$$N(\boldsymbol{z}^*) \leq N(\boldsymbol{z}) \;\text{ where } \boldsymbol{z}^* = \boldsymbol{R} \circ \boldsymbol{F}^{N(\boldsymbol{z})}(\boldsymbol{z}) \tag{25}$$

and *order preserving* condition

$$N(\boldsymbol{z}) - n \leq N(\boldsymbol{F}^n(\boldsymbol{z})) \;\text{ for any } n \tag{26}$$

The intuition behind the terminologies are explained in Appendix B along with the proof of the general VLT-CHMC algorithm.

C. Example: Rejection Avoiding HMC

Here we illustrate a use of the general VLT-CHMC framework through an algorithm of very different flavor from the special case presented in Section III A.

A step size required for stable numerical integration of Hamilton's equation (4) can vary significantly at different regions of a phase space in some application areas of HMC.[9] In such situations, the Hamiltonian may be approximately preserved along a simulated trajectory for a while until it suddenly starts to deviate wildly, leading to a proposal with little chance of acceptance. VLT-CHMC provides a way to "detect" when the trajectory becomes unstable and select an alternate state along the trajectory to transition to.

Let $\boldsymbol{F}_{\Delta t}$ be a volume-preserving and reversible map as in (5), approximating Hamiltonian dynamics. Consider a trajectory $\left\{(\boldsymbol{q}_i, \boldsymbol{p}_i) = \boldsymbol{F}^i_{\Delta t}(\boldsymbol{q}_0, \boldsymbol{p}_0)\right\}_{i=0,1,2,\ldots}$. When the trajectory becomes unstable, it can be detected by a trajectory termination criteria such as

$$b_n = \mathbb{1}\left\{ \max_{0\leq i\leq n} H(\boldsymbol{q}_i, \boldsymbol{p}_i) - \min_{0\leq i\leq n} H(\boldsymbol{q}_i, \boldsymbol{p}_i) \geq \epsilon \right\} \tag{27}$$

where $\mathbb{1}$ is an indicator function. We will actually use an alternative criteria below since this leads to a simpler algorithm implementation:

$$b_n = \mathbb{1}\Big\{ |H(\boldsymbol{q}_i, \boldsymbol{p}_i) - H(\boldsymbol{q}_{i-1}, \boldsymbol{p}_{i-1})| \geq \epsilon \quad \text{for some } i = 1, \ldots, n \Big\} \tag{28}$$

It is easy to check that the criteria (27) and (28) satisfy the properties (22) and (23) and define a valid trajectory length function N of the form (24) for Algorithm 2. We refer to the version of VLT-CHMC based on the criteria (28) as *rejection avoiding HMC*.

A proposal of rejection avoiding HMC recovers the usual HMC proposal with the trajectory length $N_{\max}$ when the fluctuation of a Hamiltonian at each step is within the error tolerance ϵ. However, upon detecting the fluctuation of magnitude larger than ϵ at the step $(\boldsymbol{q}_{i-1}, \boldsymbol{p}_{i-1}) \to (\boldsymbol{q}_i, \boldsymbol{p}_i)$, the algorithm proceeds to simulate the trajectory backward from $(\boldsymbol{q}_0, \boldsymbol{p}_0)$ and $(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*) = (\boldsymbol{q}_i, -\boldsymbol{p}_i)$ to determine the sets S and S^* according to the rule in Step 2 of Algorithm 2.

IV. NUMERICAL RESULTS

A. Geometrically tempered HMC with variable step size integrator

HMC is known to have a serious difficulty sampling from a multi-modal target density as the potential energy barriers among the modes prevents transition from one mode to another. To address this issue, Nishimura and Dunson[21] propose a version of RMHMC with a mass tensor having the property

$$|\boldsymbol{M}(\boldsymbol{q})|^{1/2} \propto \rho(\boldsymbol{q})^{1-T^{-1}} \tag{29}$$

with a *temperature* parameter $T \geq 1$. It can be shown that, with such a choice of a mass tensor, RMHMC algorithm is equivalent to the usual HMC algorithm (with a constant mass tensor) applied to a tempered distribution $\tilde{\rho}(\tilde{\boldsymbol{q}}) \propto \rho(\boldsymbol{q})^{1/T}$ on a manifold parametrized by $\tilde{\boldsymbol{q}}$. For this reason, RMHMC with the property (29) is referred to as *geometrically tempered HMC* (GTHMC) in Ref. 21.

The typical velocity of the dynamics (4) at the position $\boldsymbol{q}$ is given by the operator norm $\|\boldsymbol{M}(\boldsymbol{q})\|^{-1/2}$. This quantity, and in turn the velocity of the dynamics, necessarily becomes unboundedly large in the regions where $\rho(\boldsymbol{q})$ is small, due to the constraint (29). For this reason, the only practical way to approximate the dynamics underlying GTHMC algorithms is through a variable step size integrator with a step size proportional to $\|\boldsymbol{M}(\boldsymbol{q})\|^{1/2}$.

We take an example with a simple bimodal target density from Ref. 21. The density $\rho(\boldsymbol{q})$ is defined as a mixture of two-dimensional Gaussians with unit-variance centered at $(4, 0)$ and $(-4, 0)$. The mass tensor is chosen as

$$\begin{aligned}\boldsymbol{M}(\boldsymbol{q}) \propto\; & \rho(\boldsymbol{q})^{2\gamma\left(1-T^{-1}\right)} \boldsymbol{e}_1 \boldsymbol{e}_1^T \\ & + \rho(\boldsymbol{q})^{2(1-\gamma)(d-1)^{-1}\left(1-T^{-1}\right)} \left(\boldsymbol{I} - \boldsymbol{e}_1 \boldsymbol{e}_1^T\right)\end{aligned} \tag{30}$$

for $d^{-1} \leq \gamma \leq 1$ where $d = 2$ is the dimension of $\boldsymbol{q}$ and $\boldsymbol{e}_1 = (1, 0)$ is a standard basis vector. The mass tensors suggested in Ref. 12 and 22 have apparent resemblance to (30), but the crucial difference is that they do not satisfy (29) and consequently offer rather limited improvement over the standard HMC.

We compare the performance of CHMC and VLT-CHMC with the explicit variable step size integrator developed in Ref. 21. VLT-CHMC is run with the trajectory length function (10). The main challenge in this example to explore the phase space along the first coordinate of $\boldsymbol{q}$ due to the multi-modality along this direction. Therefore the efficiency of the sampling algorithms

Number of steps	5	10	15	20	25	30	35
Acceptance rate	0.48	0.38	0.37	0.36	0.34	0.33	0.33
ESS	75.7	180	145	83.1	103	123	101

TABLE I. ESS of CHMC along the first coordinate per 10^5 force evaluations at the various numbers of numerical integration steps. The number of steps coincides with that of force evaluations.

t	0.50	0.75	1.00	1.25	1.50	1.75	2.00
Number of steps	13	17	21	24	27	30	33
Acceptance rate	0.81	0.78	0.76	0.75	0.73	0.72	0.71
ESS	899	966	924	992	925	921	805

TABLE II. ESS of VLT-CHMC along the first coordinate per 10^5 force evaluations. The integration time t determines the trajectory lengths through the termination criteria in (10).

is summarized by the effective sample sizes (ESS) along the first coordinate of $\boldsymbol{q}$. The ESS's as well as the acceptance probabilities at different parameter settings of CHMC and VLT-CHMC are summarized in Table I and II. As predicted by our discussion in Section III A, VLT-CHMC has substantially higher acceptance probabilities and, across various parameter settings, is five times more efficient than CHMC with the optimal parameter choice. The time step size $\Delta s = .75$ for the variable step size integrator was used for all the simulations and was chosen to control the error in the Hamiltonian within a reasonable level along the trajectories. ESS's were computed using the initial monotone sequence estimator of Geyer.[23]

B. Rejection avoiding HMC

To illustrate the benefit of the rejection avoiding algorithm described in Section III C, we consider the problem of sampling from a probability density function $\rho(x, y) \propto \exp(-U(x, y))$ as plotted in Figure 1. The density $\rho(x, y)$ is constructed as a (continuous) Gaussian mixture

$$\rho(x, y) \propto \int_1^{10} \frac{1}{\sigma_\mu} \exp\left(-\frac{(x-\mu)^2}{2\,\sigma_\mu^2} - y^2\right) \mathrm{d}\mu \qquad (31)$$

where $\sigma_\mu = 0.1 + (\mu/10)^2$. The density has a property that, along the x-axis, the partial derivative $\partial_y U(x, y)$ varies substantially and so does the stable step size for the leap-frog integrator typically employed in HMC. For example, the leap-frog integrator with the step size $\Delta t \geq 0.4$ approximates the Newton's equations of motion quite accurately in the region $x > 4$, while the step size of $\Delta t \approx 0.2$ is required for a numerically stable approximation in the region $x < 2$. In in practice, such a knowledge is obviously not available to us and the appropriate step size must be determined empirically from preliminary runs of HMC. A common strategy is to pick a target acceptance rate for the HMC proposals, typically in the range $0.65 \sim 0.8$, and tune the step size accordingly.[9,24,25] This approach would suggest a step

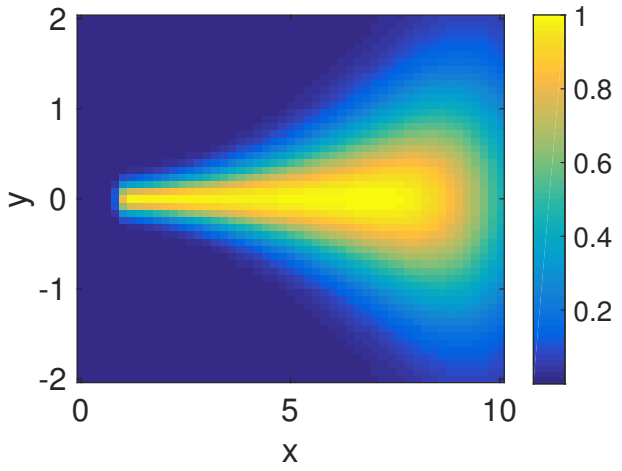

FIG. 1. A plot of (unnormalized) probability density function $\rho(x, y) \propto \exp(-U(x, y))$ used to illustrate the benefit of rejection avoiding HMC.

size well above the stability in this example, however. Figure 2 shows that the acceptance rate of HMC to be quite high even for the step size $\Delta t = 0.4$. The acceptance rate can be high despite some unstable trajectories because the region where the approximation become unstable contains relatively small, though not negligible, probability. On the other hand, the performance of HMC is severely undermined by the choice of a too large step size as can be seen in Figure 3. The ESS's for 10^6 force evaluations, estimated from ten independent simulations, are shown so that the computational cost is fixed across the experiments. The error tolerance in Hamiltonian, as in (28), for rejection avoiding HMC is set to $\epsilon = 3$. When $\Delta t = 0.2$, less than 1% of trajectories experience the error in Hamiltonian above the tolerance, so there is no practical difference between HMC with and without rejection avoidance. However, without rejection avoidance, the ESS is reduced by the factor as large as five when increasing the step size from $\Delta t = 0.2$ to $\Delta t = 0.3$. The performance degradation is less severe for rejection avoiding HMC as the algorithm concentrates the computational efforts on the stable portions of approximated trajectories.

In summary, choosing an optimal step size for HMC is difficult in practice as the choice must be made without the detailed knowledge of a target density. A step size can appear to approximate the dynamics accurately but be above the stability limit in some regions. Rejection avoiding HMC can alleviate the effect of a suboptimal step size choice and provides far more ESS's than the standard HMC in such situations.

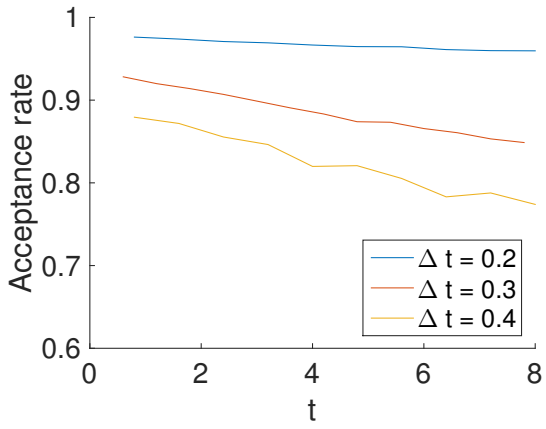

FIG. 2. Acceptance rate of HMC proposals at various settings of step size and integration time when sampling from the density shown in Figure 1.

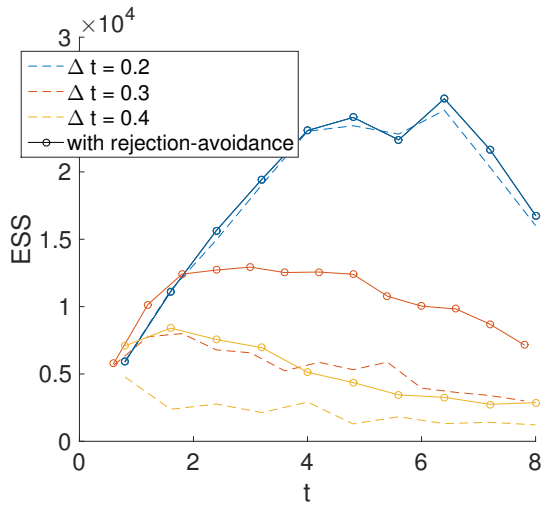

FIG. 3. ESS per 10^6 force evaluations at various settings of step size and integration time. The ESS's are for the mean estimation along the x-axis.

V. ACKNOWLEDGMENTS

We would like to thank Jiangfeng Lu for his feedback on a preliminary draft of the manuscript.

Appendix A: Derivation of limiting acceptance probability

In this section we analyse the acceptance probability of CHMC and VLT-CHMC algorithms in the special case of RMHMC with variable step size integrators as described in Section III A. We derive explicit formulas as well as useful bounds on the acceptance probabilities in the limit $\Delta s \to 0$.

1. Acceptance probability of CHMC

When approximating a time-rescaled Hamiltonian dynamics (7) with a reversible map $\boldsymbol{F}_{\Delta s}$ as in (8), the acceptance probability of the CHMC proposal from $(\boldsymbol{q}, \boldsymbol{p})$ is calculated by the formula

$$1 \wedge \frac{\rho(\boldsymbol{R} \circ \boldsymbol{F}^n_{\Delta s}(\boldsymbol{q}, \boldsymbol{p}))\, |\boldsymbol{D}\boldsymbol{F}^n_{\Delta s}(\boldsymbol{q}, \boldsymbol{p})|}{\rho(\boldsymbol{q}, \boldsymbol{p})}$$

In the limit $\Delta s \to 0$ and $n\Delta s \to s'$, the above quantity converges to

$$1 \wedge \frac{\rho(\boldsymbol{R} \circ \boldsymbol{\Phi}_{s'}(\boldsymbol{q}, \boldsymbol{p}))\, |\boldsymbol{D}\boldsymbol{\Phi}_{s'}(\boldsymbol{q}, \boldsymbol{p})|}{\rho(\boldsymbol{q}, \boldsymbol{p})}$$

where $\boldsymbol{\Phi}_s$ is the solution operator of the dynamics (7) i.e.

$$\frac{\mathrm{d}\boldsymbol{\Phi}_s}{\mathrm{d}s} = (g \circ \boldsymbol{\Phi}_s)\, \mathbf{f} \circ \boldsymbol{\Phi}_s \tag{A1}$$

where $\mathbf{f} = (\nabla_{\boldsymbol{p}} H, -\nabla_{\boldsymbol{q}} H)$. We have $\rho \circ \boldsymbol{\Phi}_{s'} = \rho$ since Hamiltonian dynamics conserves the energy and so does the time-rescaled dynamics. We also have $\rho \circ \boldsymbol{R} = \rho$, so that $\rho(\boldsymbol{R} \circ \boldsymbol{\Phi}_{s'}(\boldsymbol{q}, \boldsymbol{p})) = \rho(\boldsymbol{q}, \boldsymbol{p})$. To establish the limiting acceptance probability (9), therefore, it remains to show that $|\boldsymbol{D}\boldsymbol{\Phi}_{s'}(\boldsymbol{q}, \boldsymbol{p})| = g(\boldsymbol{\Phi}_{s'}(\boldsymbol{q}, \boldsymbol{p}))/g(\boldsymbol{q}, \boldsymbol{p})$. The Jacobian $\boldsymbol{D}\boldsymbol{\Phi}_s$ satisfies a matrix-valued differential equation $\frac{\partial}{\partial s}\boldsymbol{D}\boldsymbol{\Phi}_s = \boldsymbol{D}\mathbf{f} \circ \boldsymbol{\Phi}_s\, \boldsymbol{D}\boldsymbol{\Phi}_s$ and therefore Liouville's formula tells us that

$$|\boldsymbol{D}\boldsymbol{\Phi}_{s'}| = \exp\left(\int_0^{s'} \mathrm{tr}\,(\boldsymbol{D}\mathbf{f} \circ \boldsymbol{\Phi}_s)\, \mathrm{d}s \right)$$

A straightforward calculation shows that $\mathrm{tr}\,(\boldsymbol{D}\mathbf{f} \circ \boldsymbol{\Phi}_s) = \frac{\partial}{\partial s} \log g \circ \boldsymbol{\Phi}_s$, from which the identity $|\boldsymbol{D}\boldsymbol{\Phi}_{s'}| = g \circ \boldsymbol{\Phi}_{s'}/g$ follows.

2. Acceptance probability of VLT-CHMC

In the derivation below, we will follow the notations of Section III A 2. Namely, we set $(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*) = \boldsymbol{R} \circ \boldsymbol{F}^N(\boldsymbol{q}_0, \boldsymbol{p}_0)$, $(\boldsymbol{q}_i, \boldsymbol{p}_i) = \boldsymbol{F}^i_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0)$, and $(\boldsymbol{q}_i^*, \boldsymbol{p}_i^*) = \boldsymbol{F}^i_{\Delta s}(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*)$. The trajectory length function $N = N(t)$ is defined as in (10) and the sets S and S^* as in Algorithm 2. Note that $(\boldsymbol{q}_0, \boldsymbol{p}_0)$ is fixed, but other quantities depend on Δs, including but not limited to $(\boldsymbol{q}_i, \boldsymbol{p}_i)$'s, $N(\boldsymbol{q}_0, \boldsymbol{p}_0)$, and S. We do not denote the dependence explicitly but it is implied.

We will show that the acceptance probability of the transition from S to S^* converges to

$$1 \wedge \frac{g(\boldsymbol{\Phi}_t(\boldsymbol{q}_0, \boldsymbol{p}_0))|S^*|}{g(\boldsymbol{q}_0, \boldsymbol{p}_0)|S|} \tag{A2}$$

as $\Delta s \to 0$ while t fixed. Moreover, if $g(\boldsymbol{\Phi}_t(\boldsymbol{q}_0, \boldsymbol{p}_0)) < g(\boldsymbol{q}_0, \boldsymbol{p}_0)$, then in the limit $\Delta s \to 0$ we have $|S| = 1$ and

$$\frac{g(\boldsymbol{q}_0, \boldsymbol{p}_0)}{g(\boldsymbol{\Phi}_t(\boldsymbol{q}_0, \boldsymbol{p}_0))} - 1 \le |S^*| \le \frac{g(\boldsymbol{q}_0, \boldsymbol{p}_0)}{g(\boldsymbol{\Phi}_t(\boldsymbol{q}_0, \boldsymbol{p}_0))} + 1 \tag{A3}$$

The claimed lower bound (21) on the acceptance probability follows immediately from (A2) and (A3).

It is not difficult to show that $\mathrm{diam}(S) \to 0$ and $\mathrm{diam}(S^*) \to 0$ as $\Delta s \to 0$. This means that the elements of S (and of S^*) collapse to a single state as $\Delta s \to 0$. More precisely, for all $-\ell \le i \le r$ and $-r^* \le j \le \ell^*$,

$$\begin{aligned} \boldsymbol{F}^i_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0) &\to (\boldsymbol{q}_0, \boldsymbol{p}_0) \\ \boldsymbol{R} \circ \boldsymbol{F}^{N_0+j}_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0) &\to \boldsymbol{R} \circ \boldsymbol{\Phi}_t(\boldsymbol{q}_0, \boldsymbol{p}_0) \end{aligned} \tag{A4}$$

where $N_0 = N(\boldsymbol{q}_0, \boldsymbol{p}_0)$ and r, ℓ, r^*, ℓ^* are defined as in (13) and (15). It follows that

$$\begin{aligned} &\rho\left(\boldsymbol{F}^i_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0)\right) \left|\boldsymbol{D}\boldsymbol{F}^i_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0)\right| \to \rho(\boldsymbol{q}_0, \boldsymbol{p}_0) \\ &\rho\left(\boldsymbol{R} \circ \boldsymbol{F}^{N_0+j}_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0)\right) \left|\boldsymbol{D}\boldsymbol{F}^{N_0+j}_{\Delta s}(\boldsymbol{q}_0, \boldsymbol{p}_0)\right| \\ &\qquad \to \rho(\boldsymbol{R} \circ \boldsymbol{\Phi}_t(\boldsymbol{q}_0, \boldsymbol{p}_0)) \left|\boldsymbol{D}\boldsymbol{\Phi}_t(\boldsymbol{q}_0, \boldsymbol{p}_0))\right| \end{aligned} \tag{A5}$$

By the same argument as in Section A 1, we can show that

$$\begin{aligned} &\rho(\boldsymbol{R} \circ \boldsymbol{\Phi}_t(\boldsymbol{q}_0, \boldsymbol{p}_0)) \left|\boldsymbol{D}\boldsymbol{\Phi}_t(\boldsymbol{q}_0, \boldsymbol{p}_0))\right| \\ &\qquad = \rho(\boldsymbol{q}_0, \boldsymbol{p}_0) \frac{g(\boldsymbol{\Phi}_t(\boldsymbol{q}_0, \boldsymbol{p}_0))}{g(\boldsymbol{q}_0, \boldsymbol{p}_0)} \end{aligned} \tag{A6}$$

establishing the claimed formula (A2).

We now turn to the proof of the inequality (A3). The intuition behind the inequality and the proof below is that the size of the set $|S^*|$ is roughly equal to the number of intervals of length $\Delta s \cdot g(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*)$ that can be fit inside the interval $(t, t + \Delta s \cdot g(\boldsymbol{q}_0, \boldsymbol{p}_0))$. Denote $N_0^* = N(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*)$. By the definition of N_0^*, r^*, and ℓ^*, we must have

$$\begin{aligned} &\sum_{i=-\ell^*+1}^{N_0^*-1} \frac{\Delta s}{2} \left(g(\boldsymbol{q}^*_{i-1}, \boldsymbol{p}^*_{i-1}) + g(\boldsymbol{q}^*_i, \boldsymbol{p}^*_i)\right) \\ &\qquad < t < \sum_{i=r^*+1}^{N_0^*} \frac{\Delta s}{2} \left(g(\boldsymbol{q}^*_{i-1}, \boldsymbol{p}^*_{i-1}) + g(\boldsymbol{q}^*_i, \boldsymbol{p}^*_i)\right) \end{aligned} \tag{A7}$$

which implies that

$$\begin{aligned} &\sum_{i=-\ell^*+1}^{r^*} \frac{1}{2} \left(g(\boldsymbol{q}^*_{i-1}, \boldsymbol{p}^*_{i-1}) + g(\boldsymbol{q}^*_i, \boldsymbol{p}^*_i)\right) \\ &\qquad < \frac{1}{2}\left(g(\boldsymbol{q}^*_{N_0^*-1}, \boldsymbol{p}^*_{N_0^*-1}) + g(\boldsymbol{q}^*_{N_0^*}, \boldsymbol{p}^*_{N_0^*})\right) \end{aligned} \tag{A8}$$

Also by the definition N_0^*, r^*, and ℓ^*, we must have

$$\begin{aligned} &\sum_{i=r^*+2}^{N_0^*} \frac{\Delta s}{2} \left(g(\boldsymbol{q}^*_{i-1}, \boldsymbol{p}^*_{i-1}) + g(\boldsymbol{q}^*_i, \boldsymbol{p}^*_i)\right) \\ &\qquad < t < \sum_{i=-\ell^*}^{N_0^*-1} \frac{\Delta s}{2} \left(g(\boldsymbol{q}^*_{i-1}, \boldsymbol{p}^*_{i-1}) + g(\boldsymbol{q}^*_i, \boldsymbol{p}^*_i)\right) \end{aligned} \tag{A9}$$

which implies that

$$\begin{aligned} &\frac{1}{2}\left(g(\boldsymbol{q}^*_{N_0^*-1}, \boldsymbol{p}^*_{N_0^*-1}) + g(\boldsymbol{q}^*_{N_0^*}, \boldsymbol{p}^*_{N_0^*})\right) \\ &\qquad < \sum_{i=-\ell^*}^{r*+1} \frac{1}{2} \left(g(\boldsymbol{q}^*_{i-1}, \boldsymbol{p}^*_{i-1}) + g(\boldsymbol{q}^*_i, \boldsymbol{p}^*_i)\right) \end{aligned} \tag{A10}$$

Since $\mathrm{diam}(S) \to 0$ and $\mathrm{diam}(S^*) \to 0$ as $\Delta s \to 0$, the inequalities (A8) and (A10) converge to

$$g(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*)(|S^*| - 1) \le g(\boldsymbol{q}_0, \boldsymbol{p}_0) \le g(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*)(|S^*| + 1) \tag{A11}$$

The desired inequality (A3) is obtained by rearranging the terms in the above inequality.

Finally, we turn to the proof of the fact that $|S| \to 1$ as $\Delta s \to 0$ when $g(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*) < g(\boldsymbol{q}_0, \boldsymbol{p}_0)$. To this end, we only need to note that all the arguments in the proof of (A11) remain valid if we switch the role of $(\boldsymbol{q}_i^*, \boldsymbol{p}_i^*)$, r^*, ℓ^* and N_0^* with $(\boldsymbol{q}_i, \boldsymbol{p}_i)$, r, ℓ and N_0. This means that the inequality (A11) still holds if we switch the role of S^* with S and of $(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*)$ with $(\boldsymbol{q}_0, \boldsymbol{p}_0)$, yielding the inequality

$$g(\boldsymbol{q}_0, \boldsymbol{p}_0)(|S| - 1) \le g(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*) \le g(\boldsymbol{q}_0, \boldsymbol{p}_0)(|S| + 1)$$

In particular, we have $|S| \le \frac{g(\boldsymbol{q}_0^*, \boldsymbol{p}_0^*)}{g(\boldsymbol{q}_0, \boldsymbol{p}_0)} + 1$ and hence $|S| = 1$.

Appendix B: Justification of VLT-CHMC algorithm

As claimed in Section (III B), Algorithm 2 remains a valid algorithm when we replace the reversible map $\boldsymbol{F}_{\Delta s}$ with any reversible map and the trajectory length function N with any function of the form (24). In Section III A 3, the detailed balance condition of VLT-CHMC was derived using the notations of Algorithm 2. However, it is easy to see that the same analysis carries through when we replace the reversible map $\boldsymbol{F}_{\Delta s}$ of Section III A with any reversible map as long as the set S and S^* satisfies (12). In this section, we establish the last piece in our proof of the general VLT-CHMC algorithm; the property (12) holds whenever N satisfies the short-return (25) and order-preserving condition (26).

We consider a generic reversible map $\boldsymbol{F}$ with an associated involution $\boldsymbol{R}$ on a general phase space $\boldsymbol{z}$ as well as a generic trajectory length function N satisfying the short-return and order-preserving condition. However, all the notations and definitions directly parallel those in our presentation of the special case of VLT-CHMC in Section III A 2. Fix $\boldsymbol{z}_0$ and denote $\boldsymbol{z}_i = \boldsymbol{F}^i(\boldsymbol{z}_0)$, $\boldsymbol{z}_0^* = \boldsymbol{R} \circ \boldsymbol{F}^N(\boldsymbol{z}_0)$, and $\boldsymbol{z}_i^* = \boldsymbol{F}^i(\boldsymbol{z}_0^*)$. A trajectory function N determines the sets via the formula $S = \{\boldsymbol{z}_{-\ell}, \ldots, \boldsymbol{z}_r\}$ and $S^* = \left\{\boldsymbol{z}^*_{-\ell^*}, \ldots, \boldsymbol{z}^*_{r^*}\right\}$ where $\ell, r, \ell^*, r^* \ge 0$ are de-

fined as

$$\begin{aligned}\ell &= \max\left\{i \geq 0 : \boldsymbol{F}^N(\boldsymbol{z}_{-i}) = \boldsymbol{F}^N(\boldsymbol{z}_0)\right\}\\ r &= \max\left\{i \geq 0 : \boldsymbol{F}^N(\boldsymbol{z}_i) = \boldsymbol{F}^N(\boldsymbol{z}_0)\right\}\\ \ell^* &= \max\left\{i \geq 0 : \boldsymbol{F}^N(\boldsymbol{z}^*_{-i}) = \boldsymbol{F}^N(\boldsymbol{z}^*_0)\right\}\\ r^* &= \max\left\{i \geq 0 : \boldsymbol{F}^N(\boldsymbol{z}^*_i) = \boldsymbol{F}^N(\boldsymbol{z}^*_0)\right\}\end{aligned} \tag{B1}$$

To build the intuition behind the proof, we define a partial ordering $\preceq$ on the phase space as follows:

$$\boldsymbol{z} \preceq \tilde{\boldsymbol{z}} \quad \text{if} \quad \boldsymbol{F}^i(\boldsymbol{z}) = \tilde{\boldsymbol{z}} \text{ for } i \geq 0 \tag{B2}$$

Note that $\boldsymbol{z} \preceq \tilde{\boldsymbol{z}}$ if and only if $\boldsymbol{R}(\tilde{\boldsymbol{z}}) \preceq \boldsymbol{R}(\boldsymbol{z})$, due to the reversibility of $\boldsymbol{F}$. With this notation, the short-return condition can be expressed as

$$\boldsymbol{z} \preceq \boldsymbol{R} \circ \boldsymbol{F}^N(\boldsymbol{z}^*) \quad \text{for} \quad \boldsymbol{z}^* = \boldsymbol{R} \circ \boldsymbol{F}^N(\boldsymbol{z}) \tag{B3}$$

The condition (B3) can be interpreted intuitively as follows; according to the trajectory termination criteria imposed by N, the reverse trajectory $\boldsymbol{z}^*_0, \boldsymbol{z}^*_1, \ldots$ must terminate at $\boldsymbol{z}_0$ or at $\boldsymbol{z}_i$ for $i > 0$ before coming all the way back to $\boldsymbol{z}_0$. The order-preserving condition simply amounts to

$$\boldsymbol{F}^N(\boldsymbol{z}) \preceq \boldsymbol{F}^N(\tilde{\boldsymbol{z}}) \quad \text{if} \quad \boldsymbol{z} \preceq \tilde{\boldsymbol{z}} \tag{B4}$$

We now show how the order-preserving and short-return condition implies (12). By the order-preserving condition, we know that

$$\boldsymbol{F}^N(\boldsymbol{z}_{-\ell}) \preceq \boldsymbol{F}^N(\boldsymbol{z}_i) \preceq \boldsymbol{F}^N(\boldsymbol{z}_r) \tag{B5}$$

for all $-\ell \leq i \leq r$. On the other hand, we have $\boldsymbol{F}^N(\boldsymbol{z}_{-\ell}) = \boldsymbol{F}^N(\boldsymbol{z}_r) = \boldsymbol{R}(\boldsymbol{z}^*_0)$ by the definition of ℓ and r, so it follows that $\boldsymbol{R} \circ \boldsymbol{F}^N(\{\boldsymbol{z}_{-\ell}, \ldots, \boldsymbol{z}_r\}) = \{\boldsymbol{z}^*_0\}$.

We now turn to demonstration of $\boldsymbol{R} \circ \boldsymbol{F}^N(S^*) = \{\boldsymbol{z}_r\}$. To this end, it suffices to show $\boldsymbol{R} \circ \boldsymbol{F}^N(\boldsymbol{z}^*_0) = \boldsymbol{z}_r$ as the definition of ℓ^* and r^* combined with the order-preserving condition implies $\boldsymbol{R} \circ \boldsymbol{F}^N(\boldsymbol{z}^*_i) = \boldsymbol{R} \circ \boldsymbol{F}^N(\boldsymbol{z}^*_0)$ for all $-\ell^* \leq i \leq r^*$. Since $\boldsymbol{R} \circ \boldsymbol{F}^N(\boldsymbol{z}_r) = \boldsymbol{z}^*_0$, the short-return condition tells us $\boldsymbol{R} \circ \boldsymbol{F}^N(\boldsymbol{z}^*_0) = \boldsymbol{z}_{r+k}$ for some $k \geq 0$. To show that $k = 0$, first observe that an application of the short-return condition to the state $\boldsymbol{z}_{r+k}$ implies $\boldsymbol{z}^*_0 \preceq \boldsymbol{R} \circ \boldsymbol{F}^N(\boldsymbol{z}_{r+k})$. On the other hand, the order-preserving condition implies $\boldsymbol{F}^N(\boldsymbol{z}_r) \preceq \boldsymbol{F}^N(\boldsymbol{z}_{r+k})$ and hence $\boldsymbol{R} \circ \boldsymbol{F}^N(\boldsymbol{z}_{r+k}) \preceq \boldsymbol{R} \circ \boldsymbol{F}^N(\boldsymbol{z}_r) = \boldsymbol{z}^*_0$. The preceding inequalities together show that $\boldsymbol{R} \circ \boldsymbol{F}^N(\boldsymbol{z}_{r+k}) = \boldsymbol{z}^*_0$. Since r was defined as the largest integer i such that $\boldsymbol{R} \circ \boldsymbol{F}^N(\boldsymbol{z}_i) = \boldsymbol{z}^*_0$, it follows that $k = 0$ and $\boldsymbol{R} \circ \boldsymbol{F}^N(\boldsymbol{z}^*_0) = \boldsymbol{z}_r$.

The remaining relations in (12) as well as the fact $r^* = 0$ can be proved similarly with repeated applications of the short-return and order-preserving properties.

[1] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, "Equation of state calculations by fast computing machines," The Journal of Chemical Physics **21** (1953).

[2] G. O. Roberts and J. S. Rosenthal, "Optimal scaling for various Metropolis-hastings algorithms," Statistical Science **16**, 351–367 (2001).

[3] J. C. Mattingly, N. S. Pillai, and A. M. Stuart, "Diffusion limits of the random walk Metropolis algorithm in high dimensions," The Annals of Applied Probability **22**, 881–930 (2012).

[4] N. S. Pillai, A. M. Stuart, and A. H. Thiéry, "Optimal scaling and diffusion limits for the Langevin algorithm in high dimensions," The Annals of Applied Probability **22**, 2320–2356 (2012).

[5] S. Duane, A. Kennedy, B. J. Pendleton, and D. Roweth, "Hybrid Monte Carlo," Physics Letters B **195**, 216 – 222 (1987).

[6] B. Shahbaba, S. Lan, W. O. Johnson, and R. M. Neal, "Split Hamiltonian Monte Carlo," Statistics and Computing **24**, 339–349 (2013).

[7] A. Pakman and L. Paninski, "Auxiliary-variable exact Hamiltonian Monte Carlo samplers for binary distributions," in *Advances in Neural Information Processing Systems 26* (2013) pp. 2490–2498.

[8] A. Pakman and L. Paninski, "Exact Hamiltonian Monte Carlo for truncated multivariate Gaussians," Journal of Computational and Graphical Statistics **23**, 518–542 (2014).

[9] R. M. Neal, "MCMC using Hamiltonian dynamics," in *Handbook of Markov Chain Monte Carlo* (CRC Press).

[10] B. Leimkuhler and S. Reich, "A Metropolis adjusted nosé-hoover thermostat," ESAIM: Mathematical Modelling and Numerical Analysis **43**, 743–755 (2009).

[11] S. Lan, V. Stathopoulos, B. Shahbaba, and M. Girolami, "Markov chain Monte Carlo from Lagrangian dynamics," Journal of Computational and Graphical Statistics **24**, 357–378 (2015).

[12] Y. Fang, J. M. Sanz-Serna, and R. D. Skeel, "Compressible generalized hybrid Monte Carlo," The Journal of Chemical Physics **140** (2014), http://dx.doi.org/10.1063/1.4874000.

[13] M. Girolami and B. Calderhead, "Riemann manifold Langevin and Hamiltonian Monte Carlo methods," Journal of the Royal Statistical Society: Series B (Statistical Methodology) **73**, 123–214 (2011).

[14] M. D. Hoffman and A. Gelman, "The No-U-Turn Sampler: adaptively setting path lengths in Hamiltonian Monte Carlo," Journal of Machine Learnning Research **15**, 1593–1623 (2014).

[15] Z. Wang, S. Mohamed, and N. de Freitas, "Adaptive Hamiltonian and Riemann manifold Monte Carlo samplers," in *Proceedings of the 30th International Conference on Machine Learning*, Vol. 28 (2013) pp. 1462–1470.

[16] C. H. Bennett, "Mass tensor molecular dynamics," Journal of Computational Physics **19**, 267 – 279 (1975).

[17] M. Calvo, M. Lpez-Marcos, and J. Sanz-Serna, "Variable step implementation of geometric integrators," Applied Numerical Mathematics **28**, 1 – 16 (1998).

[18] S. Blanes and A. Iserles, "Explicit adaptive symplectic integrators for solving Hamiltonian systems," Celestial Mechanics and Dynamical Astronomy **114**, 297–317 (2012).

[19] B. Leimkuhler and S. Reich, *Simulating Hamiltonian Dynamics* (Cambridge University Press, 2005) cambridge Books Online.

[20] E. Hairer, C. Lubich, and G. Wanner, *Geometric Numerical Integration. Structure-Preserving Algorithms for Ordinary Differential Equations* (Springer-Verlag Berlin Heidelberg, 2006).

[21] A. Nishimura and D. B. Dunson, "Geometrically tempered Hamiltonian Monte Carlo," in preparation.

[22] S. Lan, J. Streets, and B. Shahbaba, "Wormhole Hamiltonian Monte Carlo," in *Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence* (2014).

[23] C. J. Geyer, "Practical markov chain monte carlo," Statist. Sci. **7**, 473–483 (1992).

[24] A. Beskos, N. Pillai, G. Roberts, J.-M. Sanz-Serna, and A. Stuart, "Optimal tuning of the hybrid Monte Carlo algorithm," Bernoulli **19**, 1501–1534 (2013).

[25] Stan Development Team, *Stan Modeling Language Users Guide and Reference Manual, Version 2.9.0* (2015).